

Next-Generation Space Internet

James Noles
Global Science and Technology, Inc.
6411 Ivy Lane, Suite 300
Greenbelt, MD 20770

Keith Scott, Mary Jo Zukoski
The MITRE Corporation
1820 Dolley Madison Blvd.
McLean, VA 22102-3481

Howard Weiss
SPARTA, Inc.
9861 Broken Land Parkway
Columbia, MD 21046

Abstract -- This paper examines elements of a communications architecture to interconnect orbiting sensors and the Internet. We examine technologies that will enable Principal Investigators (PIs) on the Internet to use familiar application interfaces to control instruments onboard spacecraft and download the large volumes of data generated by constellations of spacecraft. Our efforts are concentrated on providing the mechanisms that allow science users to request and reserve communication resources over the entire path--from their payloads to their laboratories--in a secure manner. Here we are seeking to make the most of available spacecraft communication assets, rather than attempting to provide continuous coverage of spacecraft that would not already have such connectivity.

Our work falls into four areas: 1) mechanisms to support dynamic utilization of space link communications services; 2) integrating end-to-end resource reservation mechanisms, such as the Resource reSerVation Protocol (RSVP), with the dynamic link utilization mechanisms; 3) providing user-transparency via the Mobile Internet Protocol (MobileIP) for real-time user-to-payload interaction; and 4) providing efficient end-to-end security and key management mechanisms which take advantage of existing approaches in the terrestrial environment, such as IP Security (IPSEC).

I. INTRODUCTION

Internetworking technology has radically changed the way people work and communicate on Earth. Through the Internet, users have nearly instant access to a wide range of resources including information, computational power, and data storage. Further, network-based tools for manipulating these resources can make the interface to a local database the same as to one halfway around the world.

NASA's Advanced Information Systems Technologies program is bringing these same advances to constellations of Earth-observing science instruments. The goal is to enable a highly interconnected "sensor web" of satellites that can provide long-range and detailed prediction/analysis of the Earth and its biosphere. Achieving this goal will require a large number of semi-autonomous sensors that can detect events of interest, communicate among themselves to coordinate observations, and manage the resulting large data flows. Further, there is a desire for these instruments to be accessible in "real-time" from the Internet so that they can be commanded by and deliver data to scientists and investigators who have little or no special-purpose hardware/software.

The number of spacecraft envisioned, their autonomy, and the amount of data they will generate all argue for using a shared communications infrastructure in the space segment rather than communicating with each platform individually using

some pre-determined schedule. A shared communications infrastructure, coupled with the desire to exchange information in real time with users connected to the Internet, suggests using Internet technologies such as the Internet Protocol (IP) in the space segment itself. Under this model, the orbiting sensor web becomes an extension of the Internet, where data generated on orbit may be routed through a number of satellites before reaching a downlink. Once on the ground, the data can flow across the Internet to data repositories, investigators, and other interested parties.

IP provides the ability to identify a particular endpoint (more specifically, a particular interface). In addition to this basic addressing capability, many missions will probably require other services commonly associated with the Internet suite, such as reliable data delivery and file transfers. Further, missions may also desire services that are only just now being developed and deployed within the wired Internet, such as quality of service, mobility, and security.

By extending and modifying a number of Internet technologies, we show that we can improve the communications efficiency while maintaining compatibility with the terrestrial Internet. By providing mechanisms for both ground-based controllers and onboard instruments to request and reserve system resources throughout the entire communications path, we can ensure that critical data are not lost because of network congestion. We also consider mechanisms to secure both control of the sensor web(s) and their data.

The rest of this paper is organized as follows. Section II describes current space communications capabilities as they relate to our current work. Section III describes the four areas of our work, beginning with end-to-end resource reservation and bandwidth allocation on space links. Section III then describes extensions to the MobileIP protocol to improve performance when communicating with spacecraft or other "scheduled" mobile units. Section III concludes with a discussion of security mechanisms appropriate for the space environment, and how to implement them in ways that are compatible with current efforts being deployed in the Internet. Section IV presents our conclusions and ideas for future work.

II. CURRENT CAPABILITIES

Current space missions are highly integrated. Nearly all aspects of spacecraft operations are orchestrated from the ground, so that everything from on-board resource usage to

communications are scheduled well in advance. Further, the communications of most current missions are biased towards stovepiped “application over link” architectures. Using this method, applications are intimately involved in aspects of communications that are generally associated with the data link layer, such as packet formatting and decapsulation.

The current architecture was designed to provide communications services between two known entities, and generally did not support the kinds of routing needed to support orbiting sensor webs. More recent work within the Consultative Committee for Space Data Systems (CCSDS [1]) has produced a number of data link and network layer standards that *do* support both general-purpose routing [2] and a data link layer that is better-suited to communications among multiple entities and discovery of previously unknown communications partners [3].

Bandwidth management, while an integral part of CCSDS standards from the beginning, is again highly managed. CCSDS Physical Links are logically partitioned into multiple Virtual Channels. Each Virtual Channel is allocated a portion of the total throughput of a Physical Channel that represents the total throughput of a Physical Link. CCSDS Transfer Frames are assigned to and transmitted via one of the Virtual Channels. Part of the identification of a CCSDS Transfer Frame is its Virtual Channel Identification (VCID). Packets and other application data are packaged in the Transfer Frames to be delivered across a CCSDS space link. It is this capacity of the CCSDS Transfer Frames that constitutes the bandwidth resource that can be managed by the techniques we investigate here. Figure 1 illustrates the CCSDS link architecture.

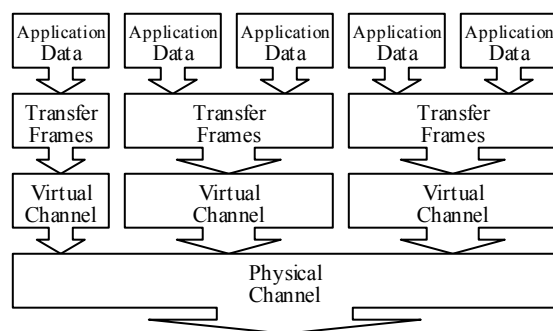


Figure 1: CCSDS link architecture.

The space science community has traditionally neglected information security. Science data was always paramount and the data unclassified. For these reasons communications were not thought of as requiring security services such as encryption and authentication. The spacecraft operators believed that it would be too difficult for a “hacker” to take control of a spacecraft because controlling one is very

difficult. As a result, most civilian space missions have been protected by the paradigm of “security through obscurity.”

III. END-TO-END COMMUNICATION: ORBITING NETWORKS TO GROUND-BASED PIS

This section describes four areas of research that we believe will help to enable the vision of large, semi-autonomous, orbiting networks communicating with Internet-based users. The goal is to allow users and orbiting instruments to interact as if they were connected via a secure, wired network while providing support to improve both manageability and performance of the communications. The ideas expressed here have been submitted to the first level of international standardization in [4, 5, 6, 7, 8].

A. End-to-End Signaling of Resource Requirements

Many aspects of spacecraft operation are carefully managed, and for good reason. Power has to be used judiciously, for example, to ensure adequate reserves for housekeeping functions. Some instruments may need to point the spacecraft in order to achieve the desired science, and competing requests must be arbitrated to prevent chaos. Treating things such as power, direction, etc. as resources allows us to bring them under the control of a resource manager. A flexible resource manager can both arbitrate among competing requests and enforce policy, such as requiring that instruments stay within envelopes specified by mission managers.

Resource reservation can also vastly improve the amount of science a mission can perform by increasing the efficiency of data transfers between space and ground. In particular, consider the transmission control protocol, TCP. TCP is the transport layer (layer-4) protocol that supports the vast majority of Internet applications today, including email (smtp), web browsing (http), file transfers (ftp), and net news (nntp). Because it provides reliable, congestion-controlled data delivery, it is likely that many of the information flows in future constellations will use TCP.

Resource reservation can improve communication efficiency by preventing congestion-based loss that would cause TCP to lower its transmission rate. Recall that TCP interprets all loss as an indication of congestion within the network, and responds by cutting its transmission rate in half. Even if all of the communications links in the space segment are managed so that there is no congestion, there is still the possibility of congestion on the ground. Thus if data from the satellite constellation does in fact flow across the Internet at large, it is possible that the locations and sources of such congestion will be completely beyond the control of mission planners and operators.

Finally, in addition to providing a unified way to allocate resources and to arbitrate competing requests, signaling of resource requirements can be coupled with mechanisms that

deliver guaranteed quality of service (QoS). Quality of service has a number of implications, and may in fact enable whole new classes of science that are not currently envisioned. On a more modest scale, because it can be used to bound message latency and jitter, QoS has the potential to simplify the design of control loops that span multiple spacecraft.

To provide the above services, we are examining the resource reservation protocol (RSVP [9]) designed by the Internet community. RSVP is an end-to-end signaling protocol that allows users to express their requirements to the network, and lets the network inform users as to whether or not those requirements can be met. When coupled with a bandwidth allocation mechanism such as that described below, RSVP can provide users with a mechanism for allocating communications resources along the entire path, from source to destination. We are also examining extensions to RSVP that will allow applications to request and reserve "local" resources such as power, pointing direction, etc.

Our results to date deal with reserving communications resources. We have simulated RSVP in a satellite constellation environment using OPNET, and have characterized the performance improvement as a result of using resource reservation. As expected, RSVP flows experienced significantly less loss than flows that did not reserve resources. Thus TCP flows using RSVP were less likely to cut their transmission rates, and maintained much higher utilization of the space link. For example, an FTP transfer using RSVP was able to complete in roughly 75 seconds. The same transfer, when congested by a non-responsive UDP flow in the ground portion of the network took nearly 1000 seconds to complete. Results also show that to adequately protect a congestion-avoiding flow (e.g. TCP) from a non-responsive congesting flow requires the allocation of both bandwidth *and* buffer space on the bottleneck link.

One problem that remains before RSVP can be effectively used between spacecraft and PI is the "multi-provider" problem. Most Internet service providers do not use RSVP to allocate bandwidth for individual flows, rather they employ it as a signaling mechanism to set up multi-protocol label-switched (MPLS) paths. Further, no provider is currently willing to allow RSVP signaling information to flow into their network from outside.

B. Bandwidth Allocation For Space Links

The RSVP protocol does not actually implement or enforce resource reservations; it merely provides the mechanism for signaling between applications and network elements. A crucial piece of our work is to allow a protocol such as RSVP to reserve resources across communications links either between spacecraft and/or between spacecraft and groundstations. This, coupled with the ability to reserve resources in the terrestrial portion of the path, will allow for

end-to-end reservation, and hence QoS. Because of the large number of missions implementing the CCSDS standards, and because of their international support, we are interested in methods for reserving bandwidth over CCSDS data link layers.

The key technical capabilities needed to allow allocation of communications resources on space links include "Traffic Classifying and Filtering", "Dynamic Route Modification" and "Output Scheduling". "Traffic Classifying and Filtering" addresses the ability to identify the critical data flows and to filter them into classes for special processing. "Dynamic Route Modification" concerns the ability to change portions of the route of the data flows without reestablishing routes end-to-end. "Output Scheduling" includes the ability to effect special processing for different classes of data vying for resources. Figure 2 illustrates these concepts at a single node in a communications path.




Figure 2: Bandwidth management concepts.

We have used the advanced traffic control capabilities of the 2.4 Linux kernel to investigate the performance improvement of bandwidth allocation by effecting "Traffic Classifying and Filtering" and "Output Scheduling." We have developed a bandwidth management simulation where multiple devices vie for limited bandwidth. The simulation examines and classifies packets based on their source addresses. Class Based Queueing (CBQ) is then used to restrict each class of

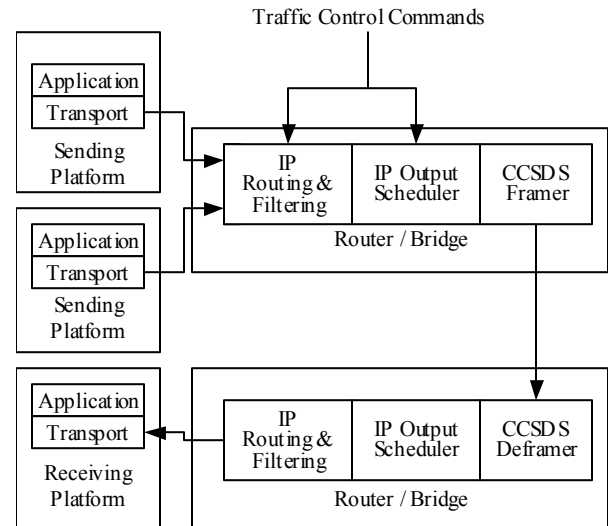


Figure 3: Experimental setup for bandwidth management.

traffic to allocated portions of the available bandwidth. Although this technique can be used over any link protocol, this demonstration has been extended to specifically manage bandwidth over CCSDS links. Using this technique, important traffic can be assured appropriate bandwidth. Coupling this capability with RSVP, which can specify the allocation parameters for the classes and links, provides a complete method for dynamically managing the bandwidth resource over CCSDS links. Figure 3 illustrates the basic topology we have investigated to date.

C. Supporting Mobility

The Internet Protocol (IP) is the glue that holds the Internet together. IP's common addressing scheme is what ensures that different computers on the Internet can find and communicate with each other, and the forwarding of IP packets by routers is the most basic service provided to end users. Thus when we speak of spacecraft that can communicate with users on the Internet, we assume the use of IP addressing.

To manage the vast number of addressable systems, IP addresses are grouped together in a topological hierarchy, and almost all forwarding is done via classless inter-domain routing, or CIDR [10]. What is important about this in our context is that *Internet routing assumes a fixed relationship between IP address and topological location within the Internet*. Low Earth Orbiting (LEO) satellites that communicate with first one ground station then another do not fit this mold. If these groundstations have different locations in the Internet topology (as they almost certainly would if run by competing providers, for example), then from the perspective of a user connected to the Internet, the satellites themselves appear to move within the network. If the IP address of the spacecraft remains the same, then it is nearly impossible to route packets to the correct groundstation in order to reach the spacecraft. Packets sent from the spacecraft to locations on the ground can generally be routed without difficulty.

Mobile nodes in the terrestrial Internet have these same problems, and researchers have developed a number of means of managing mobility from a routing perspective. One of the most popular of these is IP Mobility Support, commonly referred to as MobileIP [11]. Mobile IP specifies protocol enhancements for the transparent routing of IP datagrams to mobile nodes in the Internet. Using MobileIP, a mobile node is assigned a fixed home address, and other nodes wishing to reach the mobile use that address. If the mobile node is away from its home, a care-of address is associated with it that provides information as to its current point of attachment to the Internet. The mobile node registers the care-of address with its home agent, who then tunnels datagrams destined for the mobile agent to this care-of address. The preferred method of acquiring a care-of address is through foreign agents, in which the foreign agent acts as the endpoint of the

tunnel, un-encapsulates received datagrams, and delivers them to the mobile node.

Figure 4 illustrates the MobileIP data flows when a mobile user is connected to a foreign agent. Here the mobile has already associated itself with the foreign agent, and the foreign agent has set up an IP tunnel with the mobile's home agent. Data coming from the mobile user is routed as usual, with the mobile's home IP address (A.B.C.G in this case) as the source address. Nodes that want to send data to the mobile send it to A.B.C.G, and the data is routed toward the mobile's home agent, which intercepts the IP datagrams and tunnels them to the foreign agent. The foreign agent un-encapsulates the tunneled datagrams and forwards them to the mobile user.




Figure 4: MobileIP data flows when the mobile is associated with a foreign agent.

Mobile IP was designed to permit mobile agents to move randomly while still receiving datagrams at a fixed address. Since Mobile IP cannot predict the movement of mobile nodes, the protocol specifies several mechanisms to associate a mobile node with a mobility agent (i.e., a home or foreign agent). Spacecraft, however, do not move randomly. Contacts between spacecraft and ground stations are scheduled, with *a priori* agreement of established state. If we think of the spacecraft as a mobile node, the ground station as a foreign agent, and the control center as a home agent, then Mobile IP is directly applicable to this environment. Moreover, since the contacts are planned, the mechanisms to associate a mobile node with a locally attached mobility agent are no longer necessary. Eliminating these exchanges will free space link resources during the contact period.

We have implemented extensions to MobileIP that allow the groundstation to register on behalf of satellites, and have seen the performance increase they give through simulation. Using OPNET, we have examined a scenario with a single satellite making contact with a groundstation. As

performance measures, we considered the bandwidth saved by having the groundstation “proxy-register” the appropriate IP addresses for the satellite, as well as the time savings involved (since the extensions remove handshaking across the long-delay space link). In the future we will explore the performance benefits of MobileIP handovers using this method. We expect the benefits to be substantial given minimal signaling between spacecraft and ground to help groundstations identify acceptable handoff times (such as a beacon tone). Given this, the new groundstation will be able to “proxy register” the spacecraft’s addresses just as communications switch from the old to the new groundstation. This should ensure minimum interruption in data flow that, as mentioned in section A above, can greatly increase communications efficiency.

D. End-to-End Security for Space Missions

Providing interoperability between orbiting sensor networks and the terrestrial Internet has the potential to greatly simplify spacecraft operation. Scientists will be able to access their instruments without complicated interactions with the ground center, and data can be quickly and widely distributed throughout the Internet. Spacecraft connected to the Internet will also present an irresistible lure to hackers, opening the door to all manner of security threats, including unauthorized disclosure of data, unauthorized modification of data, and denial of service (DOS) attacks. If we are to allow commanding of spacecraft via the Internet, we will need to employ rigorous security measures to ensure that only authorized users are allowed access to the space links and the spacecraft themselves.

The Internet Engineering Task Force (IETF) Internet Protocol Security (IPSEC) working group has developed a set of security protocol standards that are just now being widely deployed in the terrestrial Internet. One drawback to using IPSEC for space missions is the additional overhead involved - a minimum of 10 bytes per IP packet. While this may not seem like much, recall that the acknowledgement stream for a TCP connection typically contains 40 or 52-byte packets, so that 10 bytes represents around 20% more overhead.

The CCSDS has developed a suite of protocols that parallel the Internet stack, but which have been extended and/or optimized for the space environment. The CCSDS security layer, Space Communications Protocol Standards - Security Protocol (SCPS-SP [12]) is a functional cousin to IPSEC, containing most of IPSEC’s capabilities but with only two bytes of overhead per IP packet. This makes it a prime candidate for use in the space segment. The reduced overhead has its price, however, and SCPS-SP is not interoperable with IPSEC. A solution to this problem is to use a trusted security gateway that can convert between IPSEC and SCPS-SP.

We have implemented a prototype “trusted gateway” that can manage IPSEC security on one side and SCPS-SP-style security on the other. Such a gateway is termed “trusted” because in order to convert between the two security protocols, the data must be momentarily “in the clear.”

A further issue complicating the use of security to protect orbiting assets is the lack of a bit-efficient key management protocol. Unlike security protocols, key management protocols generally do not add per-packet overhead. Instead they are run “out of band” periodically to distribute cryptographic information. We have investigated a number of key exchange protocols being considered for use in the Internet to determine if any could either be taken “as-is” or adapted (a la SCPS-SP) for the space environment.

Both the IPSEC and SCPS-SP protocols require the creation of *security associations*. A security association is the result of a negotiation between two parties who wish to communicate securely. Therefore, it would appear to make the most sense for the space community to either adopt the Internet Key Exchange (IKE [13]) as it presently exists, or develop a minimal profile for its use in a space communications environment while maintaining interoperability with the rest of the Internet ground infrastructure.

The most promising candidate is an operational mode of IKE termed the “aggressive exchange.” “Aggressive Exchange” allows IKE security associations, key exchanges, and authentication payloads to be transmitted together in a single IKE message. This mode reduces the number of round-trips required to establish a security association and key exchange. This is a good thing for space communications. But this reduction in overhead comes at the expense of not providing identity protection. In IKE/ISAKMP’s usual mode of operation, identities are exchanged only after a common shared secret key has been used to establish a secure communications channel. In this way the identity exchange is protected. However, when using an “aggressive exchange,” there is no shared secret in place to protect the identity exchanges. Nevertheless, the “aggressive exchange” attempts to establish all security relevant information in a single exchange. The definition of the “aggressive exchange” also allows only a single proposal and a single transform to be “negotiated” – that is, no choices are allowed.

At first blush, it would appear that the IKE/ISAKMP “aggressive exchange” is the answer to all the space community’s problems. It reduces the number of round-trips and the payload overhead required to establish a security association since it does not allow more than one proposal to be negotiated. However, its use also reduces the generality of the protocol and there is a loss of authenticated identity.

Despite the loss of authenticated identity and the ability to send multiple proposals, the security associations and key exchanges would still be interoperable with the ground-based

Internet. This means that there appears to be a way to implement an existing Internet standard in a space communications environment in a bandwidth-preserving manner, while still maintaining compatibility with the ground.

The next step is to establish a test bed to demonstrate and test that this can be done. IKE/ISAKMP servers will be set up running in an "aggressive exchange" manner. Measurements will be taken showing the overhead and latency of a non-aggressive-mode exchange versus an aggressive exchange. The differences in bandwidth utilization and round-trips will then be analyzed to determine the best approach for use of IKE/ISAKMP in the space community.

IV. CONCLUSIONS AND FUTURE WORK

We have presented a set of capabilities that can be implemented to allow communications between orbiting sensor networks and Internet-based users. These capabilities are targeted at making the sensor network appear, to the extent possible, to be directly connected to the terrestrial Internet. We note here that this real-time access model has significant drawbacks. Specifically, real-time instrument control would require investigators to keep track of details that might be unrelated to their principal mission, such as the spacecraft orbit, ground station pass times, etc. An alternate model could provide familiar interfaces without requiring knowledge of these details by using command and data caches that can be accessed via the Internet and which are synchronized with the spacecraft when appropriate. Using this model, communications between the spacecraft and the caches might use any appropriate communications technology - not necessarily the Internet suite. Future work that addresses the architectural issues associated with these and other models of interaction is sorely needed.

REFERENCES

- [1] <http://www.ccsds.org>
- [2] Space Communications Protocol Specification (SCPS) - Network Protocol (SCPS-NP), CCSDS 713.0-B-1, CCSDS, May 1999.
- [3] Proximity-1 Space Link Protocol, CCSDS 211.0-R-2, CCSDS, January 2000.
- [4] Scott, K., Noles, J., and H. Weiss, "Next Generation Space Internet Concept Document", work in progress. Available from <http://www.aist-ngsi.org>
- [5] Noles, J., "Dynamic Space Link Communications Services", work in progress. Available from <http://www.aist-ngsi.org>
- [6] Scott, K., "End-to-End Resource Provisioning for Orbiting Missions", work in progress. Available from <http://www.aist-ngsi.org>
- [7] Zukoski, M-J., "Extending IP Mobility Support to the Spacecraft Environment", work in progress. Available from <http://www.aist-ngsi.org>
- [8] Weiss, H., "End-to-End Security for Space Mission Communications", work in progress. Available from <http://www.aist-ngsi.org>
- [9] Braden, R. Ed., et. al., "Resource Reservation Protocol (RSVP) - Version 1 Functional Specification", RFC 2205, September 1997.
- [10] Fuller, V., Li, T., Yu, J., and K. Varadhan., "Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy". RFC1519, September 1993
- [11] C. Perkins, "IP Mobility Support", RFC 2002, October 1996.
- [12] Space Communications Protocol Standards - Security Protocol, CCSDS 713.5-B-1, CCSDS, May 1999.
- [13] Harkins, D. and Carrel, D., "The Internet Key Exchange (IKE)," RFC 2409, November 1998.